NORMAL FORM BACKWARD INDUCTION FOR DECISION TREES WITH COHERENT LOWER PREVISIONS

NATHAN HUNTLEY AND MATTHIAS C. M. TROFFAES

ABSTRACT. We examine normal form solutions of decision trees under typical choice functions induced by lower previsions. For large trees, finding such solutions is hard as very many strategies must be considered. In an earlier paper, we extended backward induction to arbitrary choice functions, yielding far more efficient solutions, and we identified simple necessary and sufficient conditions for this to work. In this paper, we show that backward induction works for maximality and E-admissibility, but not for interval dominance and Γ-maximin. We also show that, in some situations, a computationally cheap approximation of a choice function can be used, even if the approximation violates the conditions for backward induction; for instance, interval dominance with backward induction will yield at least all maximal normal form solutions.

1. INTRODUCTION

In classical decision theory, one aims to maximize expected utility. Such approach requires probabilities for all relevant events. However, when information and knowledge are limited, sadly, the decision maker may not be able to specify or elicit probabilities exactly. To handle this, various theories have been suggested, including lower previsions [35], which essentially amount to sets of probabilities.

In non-sequential problems, given a lower prevision, various generalizations of maximizing expected utility exist [34]. Sequential extensions of some of these alternatives have been suggested [29, 9, 13, 1, 30, 5, 14, 31], yet not systematically studied. In this paper, we study, systematically, using lower previsions, which decision criteria admit efficient solutions to sequential decision problems, by backward induction, even if probabilities are not exactly known. Our main contribution is that we prove for which criteria backward induction coincides with the usual normal form.

We study very general sequential decision problems: a subject can choose from a set of options, where each option has uncertain consequences, leading to either rewards or more options. Based on her beliefs and preferences, the subject seeks an optimal strategy. Such problems are represented by a *decision tree* [25, 18, 3].

When maximizing expected utility, one can solve a decision tree by the usual *normal form method*, or by *backward induction*. First, note that the subject can specify, in advance, her actions in all eventualities. In the normal form, she simply chooses a specification which maximizes her expected utility. However, in larger problems, the number of specifications is gargantuan, and the normal form is not feasible.

Fortunately, backward induction is far more efficient. We find the expected utility at the final decision nodes, and then replace these nodes with the maximum expected utility. The previously penultimate decision nodes are now ultimate, and the process repeats until the root is

reached. Backward induction is guaranteed to coincide with the normal form [25] if probabilities are non-zero [8, p. 44].

The usual normal form method works easily with decision criteria for lower previsions: apply it to the set of all strategies. Generalizing backward induction is harder, as no single expectation summarizes all relevant information about substrategies, unlike with expected utility. We follow Kikuti et al. [14], and instead replace nodes with sets of optimal substrategies, moving from right to left in the tree, eliminating strategies as we go. De Cooman and Troffaes [5] presented a similar idea for dynamic programming.

In this general setting, normal form and backward induction can differ, as noted by many [29, 20, 30, 14, 5, 13, 1]. However, for some decision criteria the methods always coincide. In [11], we found conditions for coincidence. In this paper, we expand the work begun in [12], and investigate what works for lower previsions, finding that *maximality* and *E-admissibility* work, but the others do not.

This coincidence is of interest for at least two reasons. First, as mentioned, the normal form is not feasible for larger trees, whereas backward induction can eliminate many strategies early on, hence being far more efficient. Secondly, one might argue that a solution where the two methods differ is philosophically flawed [11, 8, 29, 21].

The paper is organized as follows. Section 2 explains decision trees and introduces notation. Section 3 presents lower previsions and their decision criteria, and demonstrates normal form backward induction on a simple example. Section 4 formally defines the two methods, and characterizes their equivalence, which is applied in Section 5 to lower previsions. Section 6 discusses a larger example. Section 7 concludes. Readers familiar with decision trees and lower previsions can start with Sections 3.3 and 6.

2. Decision Trees

2.1. **Definition and Example.** Informally, a decision tree [18, 3] is a graphical causal representation of decisions, events, and rewards. Decision trees consist of a rooted tree [7, p. 92, Sec. 3.2] of decision nodes, chance nodes, and reward leaves, growing from left to right. The left hand side corresponds to what happens first, and the right side to what happens last.

Consider the following example. Tomorrow, a subject is going for a walk in the lake district. It may rain (E_1), or not (E_2). The subject can either take a waterproof (d_1), or not (d_2). But the subject may also choose to buy today's newspaper, at cost c, to learn about tomorrow's weather forecast (d_S), or not ($d_{\overline{S}}$), before leaving for the lake district. The forecast has two possible outcomes: predicting rain (S_1), or not (S_2).

The corresponding decision tree is depicted in Figure 1. Decision nodes are depicted by squares, and chance nodes by circles. From each node, a number of branches emerge, representing decisions at decision nodes and events at chance nodes. The events from a node form a partition of the possibility space: exactly one of the events will take place. Each path in a decision tree corresponds to a sequence of decisions and events. The reward from each such sequence appears at the right hand end of the branch.

2.2. **Notation.** A particular decision tree can be seen as a combination of smaller decision trees: for example, one could draw the subtree corresponding to buying the newspaper, and also draw the subtree corresponding to making an immediate decision. The decision tree for the full problem is then formed by joining these two subtrees at a decision node.

So, we can represent a decision tree by its subtrees and the type of its root node. Let $T_1, \ldots, T_n$ be decision trees. If T combines the trees at a decision node, we write
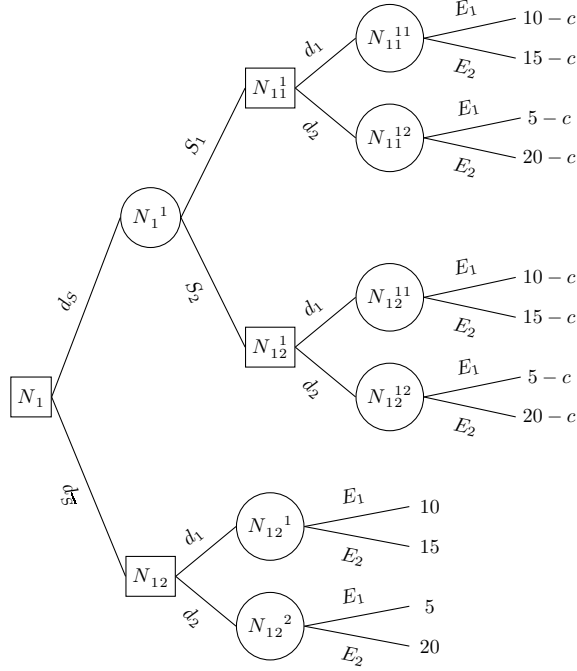
$$T = \bigsqcup_{i=1}^{n} T_i.$$

FIGURE 1. A decision tree for walking in the lake district.

If T combines the trees at a chance node, with subtree T_i being connected by event E_i ($E_1, \ldots, E_n$ is a partition of the possibility space) we write

$$T = \bigodot_{i=1}^{n} E_i T_i.$$

For instance, for the tree of Fig. 1 with $c = 1$, we write

$$(S_1(T_1 \sqcup T_2) \odot S_2(T_1 \sqcup T_2)) \sqcup (U_1 \sqcup U_2)$$

with, where we denoted the reward nodes by their utility,

$$T_1 = E_1 9 \odot E_2 14 \qquad\qquad U_1 = E_1 10 \odot E_2 15$$
$$T_2 = E_1 4 \odot E_2 19 \qquad\qquad U_2 = E_1 5 \odot E_2 20$$

The above notation shall prove very useful when considering recursive definitions.

In this paper we often consider subtrees of larger trees. For subtrees, it is important to know the events that were observed in the past. Two subtrees with the same configuration of nodes and arcs may have different preceding events, and should be treated differently. Therefore we associate with every decision tree T an event $\mathrm{ev}(T)$ representing the intersection of all the events on chance arcs that have preceded T.

Definition 1. *A subtree of a tree T obtained by removal of all non-descendants of a particular node N is called* the subtree of T at N *and is denoted by* $\mathrm{st}_N(T)$.

These subtrees are called 'continuation trees' by Hammond [8].

Consider all possible ways that sets of decision trees $\mathcal{T}_1, \ldots, \mathcal{T}_n$ can be combined. Our notation easily extends. For any partition $E_1, \ldots, E_n$,

$$\bigodot_{i=1}^{n} E_i \mathcal{T}_i = \left\{ \bigodot_{i=1}^{n} E_i T_i : T_i \in \mathcal{T}_i \right\}.$$

For any sets of consistent decision trees $\mathcal{T}_1, \ldots, \mathcal{T}_n$,

$$\bigsqcup_{i=1}^{n} \mathcal{T}_i = \left\{ \bigsqcup_{i=1}^{n} T_i : T_i \in \mathcal{T}_i \right\}.$$

For convenience we only work with decision trees for which there is no event arc that is impossible given preceding events.

Definition 2. *A decision tree T is called* consistent *if for every node N of T,*

$$\operatorname{ev}(\operatorname{st}_N(T)) \neq \emptyset.$$

Clearly, if a decision tree T is consistent, then for any node N in T, $\operatorname{st}_N(T)$ is also consistent. Considering only consistent trees is not really a restriction, since inconsistent trees would only be drawn due to an oversight and could easily be made consistent.

2.3. **Solving Decision Trees with Probabilities and Utilities.** We give a brief overview of the standard method of solving a decision tree when probabilities of events are known. Suppose in Fig. 1 we have $p(S_1) = 0.6$, $p(E_1|S_1) = 0.7$, and $p(E_1|S_2) = 0.2$, so $p(E_1) = 0.5$. We first calculate the expected utility of the final chance nodes. For example, the expected utility at N_{11}^{11} is $0.7(10 - c) + 0.3(15 - c) = 11.5 - c$, and the expected utility at N_{11}^{12} is $9.5 - c$.

We now see that at N_{11}^{1} it is better to choose decision d_1. We then *replace* N_{11}^{1} and its subtree with the expected utility of N_{11}^{11}: $11.5 - c$. Also follow this procedure for N_{12}^{1} and N_{12}, and the tree has been reduced by a stage. We find that d_2 is optimal at N_{12}^{1} with value $17 - c$, and at N_{12} both decisions are optimal with value 12.5.

Next, take expected utility at N_1^{1}, which is $0.6(11.5 - c) + 0.4(17 - c) = 13.7 - c$. At N_1, we therefore take decision d_S if $c \leq 1.2$ and $d_{\overline{S}}$ if $c \geq 1.2$. This procedure is illustrated in Fig. 2, where the dashed lines indicate decision arcs that are rejected because their expected utility is too low (for any specific $c \neq 1.2$, either d_S or $d_{\overline{S}}$ would be dashed).

This method can only be carried out if the subject has assessed precise probabilities and utilities and wishes to maximize expected utility. It may be that the subject is unable or unwilling to comply with these requirements. The next section considers a possible solution, and demonstrates how backward induction can be generalized.

3. CREDAL SETS AND LOWER PREVISIONS

First, we outline a straightforward generalization of the theory of probability, allowing the subject to model uncertainty in cases where too little information is available to identify a unique probability distribution (see for instance [2, 26, 6, 33, 36, 37, 35]).

3.1. **Gambles, Credal Sets and Lower Previsions.** The *possibility space* Ω is the set of all possible states of the world. Elements of Ω are denoted by ω. Subsets of Ω are called *events*, and are denoted by capital letters A, B, etc. The arcs emerging from chance nodes in a decision tree correspond to events.

A *gamble* is a function $X \colon \Omega \to \mathbb{R}$, interpreted as an uncertain reward: should $\omega \in \Omega$ be the true state of the world, the gamble X will yield the reward $X(\omega)$.

A *probability mass function* is a non-negative real-valued function $p \colon \Omega \to \mathbb{R}^+$ whose values sum to one [28, p. 138, Sec. 4.2]. For convenience, we will write $p(A)$ for $\sum_{\omega \in A} p(\omega)$.

FIGURE 2. Solving the Lake District problem with expected utility.

For the purpose of this paper, we assume that

- Ω is finite,
- rewards are expressed in utiles,
- the subject can express her beliefs by means of a closed convex set $\mathcal{M}$ of probability mass functions p ($\mathcal{M}$ is called the *credal set*), and
- each probability mass function $p \in \mathcal{M}$ satisfies $p(\omega) > 0$, for all $\omega \in \Omega$.

Under the above assumptions, each p in $\mathcal{M}$ determines a conditional expectation

$$E_p(X|A) = \frac{\sum_{\omega \in A} X(\omega)p(\omega)}{p(A)},$$

and the whole set $\mathcal{M}$ determines a conditional lower and upper expectation

$$\underline{P}(X|A) = \min_{p \in \mathcal{M}} E_p(X|A) \qquad\qquad \overline{P}(X|A) = \max_{p \in \mathcal{M}} E_p(X|A),$$

and this for every gamble X and every non-empty event A.

The functional $\underline{P}$ is called a *coherent conditional lower prevision*, and similarly, $\overline{P}$ is called a *coherent conditional upper prevision*. Although here we have defined these by means of a set of probability measures, there are different ways of obtaining and interpreting lower and upper previsions (see for instance Miranda [22] for a survey).

Following De Finetti [6], where it is convenient we denote the indicator gamble

$$I_A(\omega) = \begin{cases} 1 & \text{if } \omega \in A \\ 0 & \text{if } \omega \notin A \end{cases}$$

of an event A also simply by A: for instance, if X is a gamble, and A is an event, then AX is just a shorthand notation for $I_A X$.

Below are some properties of coherent conditional lower and upper previsions that we require later (see Williams [36, 37] or Walley [35] for proofs).

Proposition 3. *For all non-empty events A, B, gambles X, Y, and constants $\lambda > 0$:*

  (i) *If $AX = AY$ then $\underline{P}(X|A) = \underline{P}(Y|A)$.*
  (ii) $\underline{P}(X|A) + \underline{P}(Y|A) \leq \underline{P}(X + Y|A) \leq \underline{P}(X|A) + \overline{P}(Y|A)$.
  (iii) $\underline{P}(\lambda X|A) = \lambda \underline{P}(X|A)$ *and* $\overline{P}(\lambda X|A) = \lambda \overline{P}(X|A)$.
  (iv) $\underline{P}(X|A) = -\overline{P}(-X|A)$.
  (v) $\underline{P}(A(X - \underline{P}(X|A \cap B))|B) = 0$.

Property (v) generalizes the *generalized Bayes rule*, which describes how conditional lower previsions are linked to unconditional lower previsions. For instance, if $\underline{P}(A|B) > 0$, then $\underline{P}(\cdot|A \cap B)$ is uniquely determined by $\underline{P}(\cdot|B)$ via property (v) [35, p. 297, Thm. 6.4.1].

3.2. **Choice Functions and Optimality.** Suppose a subject must choose a gamble from a set $\mathcal{X}$. In classical decision theory, a gamble is optimal if its expectation is maximal. More generally, given a credal set $\mathcal{M}$, the subject might consider as optimal for example any gambles whose expectation is maximal for at least one $p \in \mathcal{M}$: in fact, most criteria determine optimal decisions by comparison of gambles. So, we suppose that the subject has a way of determining an optimal subset of any set of gambles, conditional on any event A (think of A as $\mathrm{ev}(T)$):

Definition 4. *A* choice function opt *is an operator that, for each non-empty event A, maps each non-empty finite set $\mathcal{X}$ of gambles to a non-empty subset of this set:*

$$\emptyset \neq \mathrm{opt}(\mathcal{X}|A) \subseteq \mathcal{X}.$$

Note that common uses of choice functions in social choice theory, such as by Sen [32, p. 63, ll. 19–21], do not consider conditioning on events, and define choice functions for arbitrary sets of options, rather than for gambles only.

The interpretation of a choice function is that when the subject can choose among the elements of $\mathcal{X}$, having observed A, she would only choose from $\mathrm{opt}(\mathcal{X}|A)$. Therefore, we say the elements of $\mathrm{opt}(\mathcal{X}|A)$ are *optimal* (relative to $\mathcal{X}$ and A). Note that the subject may not consider them equivalent: adding a small incentive to choose a particular optimal option would not necessarily make it the single preferred option.

We now consider four popular choice functions that have been proposed for choosing between gambles given a coherent lower prevision. Further discussion of the criteria presented here can be found in Troffaes [34].

3.2.1. *Maximality.* Maximality is based on the following strict partial preference order $>_{\underline{P}|A}$.

Definition 5. *Given a coherent lower prevision $\underline{P}$, for any two gambles X and Y we write $X >_{\underline{P}|A} Y$ whenever $\underline{P}(X - Y|A) > 0$.*

The partial order $>_{\underline{P}|A}$ gives rise to the choice function *maximality*, proposed by Condorcet [4, pp. lvj–lxix, 4.^{e} Exemple], Sen [32], and Walley [35], among others.[1]

Definition 6. *For any non-empty finite set of gambles $\mathcal{X}$ and each event $A \neq \emptyset$,*

$$\mathrm{opt}_{>_{\underline{P}|\cdot}}(\mathcal{X}|A) = \{X \in \mathcal{X} : (\forall Y \in \mathcal{X})(Y \not>_{\underline{P}|A} X)\}.$$

---

[1]Because all probabilities in $\mathcal{M}$ are assumed to be strictly positive, Walley's admissibility condition is implied and hence omitted in Definition 6.

3.2.2. *E-admissibility.* Another criterion is E-admissibility, proposed by Levi [17]. Recall that $\underline{P}(\cdot|A)$ is the lower envelope of $\mathcal{M}$. For each $p \in \mathcal{M}$ we can maximize expected utility:

$$\text{opt}_p(\mathcal{X}|A) = \{X \in \mathcal{X} \colon (\forall Y \in \mathcal{X})(E_p(Y|A) \leq E_p(X|A))\}.$$

Then the set of E-admissible options is defined by:

Definition 7. *For any non-empty finite set of gambles $\mathcal{X}$ and each event $A \neq \emptyset$,*

$$\text{opt}_{\mathcal{M}}(\mathcal{X}|A) = \bigcup_{p \in \mathcal{M}} \text{opt}_p(\mathcal{X}|A).$$

A gamble X is therefore E-admissible when it maximizes expected utility under at least one $p \in \mathcal{M}$. Any E-admissible gamble is maximal [35, p. 162, ll. 26–28].

3.2.3. *Interval Dominance.* Interval dominance is based on the strict partial preference order $\sqsupset_{\underline{P}|A}$.

Definition 8. *Given a coherent lower prevision $\underline{P}$, for any non-empty event A and any two A-consistent gambles X and Y we write $X \sqsupset_{\underline{P}|A} Y$ whenever $\underline{P}(X|A) > \overline{P}(Y|A)$.*

This ordering induces a choice function usually called *interval dominance* [38, 34]:

Definition 9. *For any non-empty finite set of gambles $\mathcal{X}$ and each event $A \neq \emptyset$,*

$$\text{opt}_{\sqsupset_{\underline{P}|}}(\mathcal{X}|A) = \{X \in \mathcal{X} \colon (\forall Y \in \mathcal{X})(Y \not\sqsupset_{\underline{P}|A} X)\}.$$

The above criterion was apparently first introduced by Kyburg [15] and was originally called stochastic dominance.

3.2.4. *Γ-maximin.* Γ-maximin selects gambles that maximizes the minimum expected reward.

Definition 10. *For any non-empty finite set of gambles $\mathcal{X}$ and each event $A \neq \emptyset$,*

$$\text{opt}_{\underline{P}}(\mathcal{X}|A) = \{X \in \mathcal{X} \colon (\forall Y \in \mathcal{X})(\underline{P}(X|A) \geq \underline{P}(Y|A))\}.$$

Γ-maximin is a total preorder, and so usually selects a single gamble regardless of the degree of uncertainty in $\underline{P}$. Γ-maximin can be criticized for being too conservative (see Walley [35, p. 164]), as it only takes into account the worst possible scenario.

3.3. **Sequential Problems Using Lower Previsions.** Consider again the lake district problem depicted in Fig. 1, but now suppose that the subject has specified a coherent lower prevision, instead of a singe probability measure. For this example, we consider an ϵ-contamination model: with probability $1 - \epsilon$, observations follow a given probability mass function p, and with probability ϵ, observations follow an unknown arbitrary distribution. One can easily check that, under this model, the lower expectation for a gamble X is

$$\underline{P}(X) = (1 - \epsilon)E_p(X) + \epsilon \inf X$$

The conditional lower expectation is [35, p. 309]

$$\underline{P}(X|A) = \frac{(1 - \epsilon)E_p(AX) + \epsilon \inf_{\omega \in A} X(\omega)}{(1 - \epsilon)p(A) + \epsilon}$$

As before, let $p(S_1) = 0.6$, $p(E_1|S_1) = 0.7$, and $p(E_1|S_2) = 0.2$, so $p(E_1) = 0.5$. Let $\epsilon = 0.1$.

Naively, she could solve the problem with the usual *normal form method*: she lists all possible strategies (actions to take in all eventualities), finds the corresponding gambles, and applies a suitable choice function, say maximality.

Table 1 lists all strategies and their gambles. Each strategy gives a reward determined entirely by ω, and hence has a corresponding gamble. For example, the gamble for the last strategy is

$$(5 - c)S_1 E_1 + (20 - c)S_1 E_2 + (10 - c)S_2 E_1 + (15 - c)S_2 E_2 = S_1 Y + S_2 X - c,$$

strategy	gamble
$d_{\overline{S}}$, then d_1	X
$d_{\overline{S}}$, then d_2	Y
d_S, then d_1 if S_1 and d_1 if S_2	$X - c$
d_S, then d_2 if S_1 and d_2 if S_2	$Y - c$
d_S, then d_1 if S_1 and d_2 if S_2	$S_1 X + S_2 Y - c$
d_S, then d_2 if S_1 and d_1 if S_2	$S_1 Y + S_2 X - c$

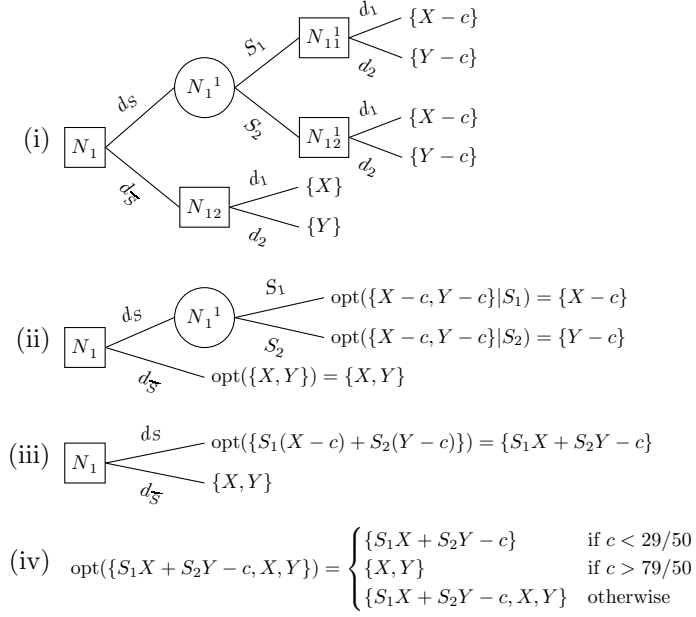TABLE 1. Strategies and gambles for the lake district problem.

FIGURE 3. Solving the lake district example by normal form backward induction.

with $X = 10E_1 + 15E_2$ and $Y = 5E_1 + 20E_2$. Recall that $(5-c)S_1 E_1$ is just a shorthand notation for $(5-c)I_{S_1} I_{E_1}$, and similarly for all other terms.

Maximality can then be applied to find the optimal gambles: this requires comparison of all six gambles at once. Skipping the details of this calculation, for instance with $c = 0.5$, we find that we should buy the newspaper and follow its advice.

However, could we think of a backward induction scheme which might not require comparison of all six gambles at once? Obviously, any such scheme will not work as easily as in Section 2.3, because we are not maximizing for a single number (i.e. expected utility). Instead, *we retain the optimal strategies in subtrees*, as illustrated in Fig. 3.

Denote subtrees at a particular node N_*^* by $T_*^* = \mathrm{st}_{N_*^*}(T)$.

(i) First, write down the gambles at the final chance nodes. For example, at N_{11}^{11} the gamble is $(10 - c)E_1 + (15 - c)E_2 = X - c$, and similarly for all others.

(ii) Let us first deal with the branch corresponding to refusing the newspaper. At the decision node N_{12}, we have a choice between two strategies that correspond to the gambles X and Y. We also have $\mathrm{ev}(T_{12}) = \Omega$. So to determine the optimal strategies in this subtree,

we must compare these two gambles unconditionally:

$$\underline{P}(X - Y) = \underline{P}(Y - X) = -5\epsilon = -1/2,$$

so at T_{12} the strategies d_1 and d_2 are both optimal.

Now we move to the branch corresponding to buying the newspaper. At N_{11}^1, we need to compare $X - c$ and $Y - c$. We have that $\mathrm{ev}(T_{11}^1) = S_1$, and

$$\underline{P}((X - c) - (Y - c)|S_1) = \tfrac{6-31\epsilon}{3+2\epsilon} > 0,$$

so $X - c >_{\underline{P}|S_1} Y - c$ and the uniquely optimal strategy is d_1. Next, considering N_{12}^1, we see that $\mathrm{ev}(T_{12}^1) = S_2$, and

$$\underline{P}((Y - c) - (X - c)|S_2) = \tfrac{6-31\epsilon}{2+3\epsilon} > 0,$$

so the optimal strategy here is d_2.

(iii) Moving to N_1^1, we see that only one of the original four strategies remains: "d_1 if S_1 and d_2 if S_2", corresponding to the gamble $S_1 X + S_2 Y - c$.

(iv) Finally, considering the entire tree T, three strategies are left: "d_1 if S_1 and d_2 if S_2"; "$d_{\overline{S}}$, then d_1"; "$d_{\overline{S}}$, then d_2". Therefore we need to find

$$\mathrm{opt}_{>_{\underline{P}|\cdot}}(\{S_1 X + S_2 Y - c, X, Y\}).$$

We have

$$\underline{P}(X - (S_1 X + S_2 Y - c)) = c - (6 + 19\epsilon)/5 = c - 79/50$$
$$\underline{P}((S_1 X + S_2 Y - c) - X) = (6 - 31\epsilon)/5 - c = 29/50 - c$$
$$\underline{P}(Y - (S_1 X + S_2 Y - c)) = c - (6 + 19\epsilon)/5 = c - 79/50$$
$$\underline{P}((S_1 X + S_2 Y - c) - Y) = (6 - 31\epsilon)/5 - c = 29/50 - c$$

Concluding (see Fig. 3(iv)):
- if the newspaper costs less than 29/50, we should buy and follow its advice.
- if it costs more than 79/50, we do not buy, but have insufficient information to decide whether to take the waterproof or not.
- if the newspaper costs between 29/50 and 79/50, we can take any of the three remaining options.

Comparing this with the solution calculated in Section 2.3, we observe that the imprecision has created a range of c for which it is unclear whether buying the newspaper is better than not, rather than the single value for c in the precise case. Despite this, should the subject decide to buy the newspaper, she will follow the same policy in both cases: take the waterproof only if the newspaper predicts rain. Finally it should be noted that, although in both cases both $d_{\overline{S}}d_1$ and $d_{\overline{S}}d_2$ are involved in optimal normal form decisions for some values of c, in the precise case this is because they are *equivalent* and in the imprecise case they are *incomparable*. A tiny increase in value for, say, not taking the waterproof and no rain, would make $d_{\overline{S}}d_1$ always non-optimal under E_p but still optimal under $\underline{P}$ for $c \geq 29/50$.

For this particular example, it is easy though tedious to check that backward induction gives the same answer as the usual normal form method (that is, applying maximality to all six gambles), for any value of c. It is easy to find choice functions and decision trees where this does not work [16, 13, 30]. We want to know for which of the coherent lower previsions choice functions the two methods agree. To answer this, we invoke a theorem relating to general choice functions, outlined in the next section.

4. Normal Form Solutions for Decision Trees

We now introduce the necessary terminology for examining our two methods of solution in detail, and provide theorems stating when they coincide. These two methods yield *normal form solutions*, that is, sets of optimal strategies at the root node.

4.1. **Normal Form Decisions, Solutions, Operators, and Gambles.** Suppose the subject specifies a decision for each eventuality, and resolves to follow this policy. She now has no choice at decision nodes, and her reward is entirely determined by the state of nature. This corresponds to a reduced decision tree obtained by taking the initial tree and removing all but one of the arcs at each decision node. Such a reduced tree is called a *normal form decision*, and represents what we called a "strategy" in Section 3.3. We denote the set of all normal form decisions of T by $\mathrm{nfd}(T)$.

It is unlikely that the subject can specify a single optimal normal form decision for all problems. Nevertheless, she might be able to eliminate some unacceptable ones: a *normal form solution* of a decision tree T is simply a non-empty subset of $\mathrm{nfd}(T)$. A *normal form operator* is then a function mapping every decision tree to a normal form solution of that tree. The two methods we investigate are normal form operators.

As we saw in Section 3.3, the reward for a normal form decision is determined entirely by the events that take place. That is, a normal form decision has a corresponding gamble, which we call a *normal form gamble*. The set of all normal form gambles associated with a decision tree T is denoted by $\mathrm{gamb}(T)$, so gamb is an operator on trees that yields the set of all gambles induced by normal form decisions of the tree.

We will need to know when a set of gambles can be represented by a consistent decision tree (as defined earlier in Section 2.2):

Definition 11. *Let A be any non-empty event, and let $\mathcal{X}$ be a set of gambles. Then the following conditions are equivalent; if they are satisfied, we say that $\mathcal{X}$ is A-consistent.*

(A) *There is a consistent decision tree T with $\mathrm{ev}(T) = A$ and $\mathrm{gamb}(T) = \mathcal{X}$.*
(B) *For all $r \in \mathbb{R}$ and all $X \in \mathcal{X}$ such that $X^{-1}(r) \neq \emptyset$, it holds that $X^{-1}(r) \cap A \neq \emptyset$.*

Proof of equivalence of (A) and (B) is fairly straightforward, whence omitted here.
The following notation proves convenient for normal form gambles at chance nodes.

Definition 12. *For any events E_1, ..., E_n which form a partition, and any finite family of sets of gambles $\mathcal{X}_1$, ..., $\mathcal{X}_n$, we define the following set of gambles:*

$$(1) \qquad \sum_{i=1}^{n} E_i \mathcal{X}_i = \left\{ \sum_{i=1}^{n} E_i X_i \colon X_i \in \mathcal{X}_i \right\}$$

4.2. **Normal Form Operator Induced by a Choice Function.** We can now formalize the simple normal form method described at the start of Section 3.3. Listing all strategies corresponds to finding $\mathrm{nfd}(T)$. Listing their corresponding gambles corresponds to finding $\mathrm{gamb}(T)$. Then calculate $\mathrm{opt}(\mathrm{gamb}(T)|\mathrm{ev}(T))$, and find all elements of $\mathrm{nfd}(T)$ that induced these optimal gambles. The solution is then the set of all these normal form decisions. Formally,

Definition 13. *Given any choice function* opt, *and any decision tree T with $\mathrm{ev}(T) \neq \emptyset$,*

$$\mathrm{norm}_{\mathrm{opt}}(T) = \{U \in \mathrm{nfd}(T) \colon \mathrm{gamb}(U) \subseteq \mathrm{opt}(\mathrm{gamb}(T)|\mathrm{ev}(T))\}.$$

The following important equality follows immediately:

$$(2) \qquad \mathrm{gamb}(\mathrm{norm}_{\mathrm{opt}}(T)) = \mathrm{opt}(\mathrm{gamb}(T)|\mathrm{ev}(T)).$$

Let us demonstrate this definition on lake district problem. For any particular strategy U, say for instance "buy the newspaper and take the waterproof only if the newspaper predicts rain", we can calculate its associated gamble U, which is in our instance

$$X = (10 - c)E_1S_1 \oplus (15 - c)E_2S_1 \oplus (5 - c)E_1S_2 \oplus (20 - c)E_2S_2.$$

We check whether X is optimal in the set of all gambles associated with T, that is, whether $X \in \mathrm{opt}(\mathrm{gamb}(T)|\mathrm{ev}(T))$. If so, then $\mathrm{gamb}(U) = \{X\} \subseteq \mathrm{opt}(\mathrm{gamb}(T)|\mathrm{ev}(T))$ and so $U \in \mathrm{norm}_{\mathrm{opt}}(T)$. Otherwise, $\mathrm{gamb}(U) = \{X\} \nsubseteq \mathrm{opt}(\mathrm{gamb}(T)|\mathrm{ev}(T))$ and so $U \notin \mathrm{norm}_{\mathrm{opt}}(T)$. This procedure for each strategy in T will determine $\mathrm{norm}_{\mathrm{opt}}(T)$.

4.3. **Normal Form Backward Induction.** The operator $\mathrm{norm}_{\mathrm{opt}}$ is a natural and popular choice, but for practical or philosophical reasons one may wish to be able to find it by backward induction. Even for an almost trivial problem such as Fig. 1 there are already six normal form decisions. If a tree T has at least n decision nodes in every path from the root to any leaf, and each decision node has at least two children, then there will be at least 2^n normal form decisions associated with T (and often a lot more). Working with sets of 2^n gambles may be impractical for large n, particularly for maximality and E-admissibility, so a method that may avoid applying the choice function on a large set is necessary.

Implementation of backward induction is easy when there is a unique choice at every node, but a choice function may not have this property, so we need to adapt the traditional approach. The technique informally introduced in Section 3.3 is a generalization of the method of Kikuti et al. [14], where the only difference is that we apply our choice function at all nodes, not just decision nodes. Although the focus of Kikuti et al. is also on uncertainty models represented by coherent lower previsions, their approach, and so our generalization, can be used for any choice function on gambles.

The goal of our backward induction algorithm is to reach a normal form solution of T by finding normal form solutions of subtrees of T, and using these to remove some elements of $\mathrm{nfd}(T)$ before applying opt. A formal definition requires many definitions that hinder clarity, so in this paper we prefer a more intuitive informal approach. Rigorous treatment of our backward induction operator can be found in [10, 11].

The algorithm moves from right to left in the tree as follows. At a subtree $\mathrm{st}_N(T)$, find the set of normal form decisions $\mathrm{nfd}(T)$, but remove any strategies that contain substrategies judged non-optimal at any descendent node. For example, in Fig. 1, at $N_1{}^1$ there are four normal form decisions, but in the example the substrategy d_2 was removed at N_{11}^1, and d_1 was removed at N_{12}^1, so the only strategy at $N_1{}^1$ that is retained is d_1 if S_1, d_2 if S_2. Next, find the corresponding gambles of all surviving normal form decisions, apply opt, and transform back to optimal normal form decisions. Move to the next layer of nodes, and continue until the root node is reached. This yields a set of normal form decisions at the root node, that is, the algorithm corresponds to a normal form operator, which we call $\mathrm{back}_{\mathrm{opt}}$.

A further example using maximality, but this time using the notation of this section, can be found in Section 6, if further clarification is required.

In [10, 11], we found four necessary and sufficient properties on opt for $\mathrm{back}_{\mathrm{opt}}$ and $\mathrm{norm}_{\mathrm{opt}}$ to coincide for any consistent decision tree.

Property 1 (Backward conditioning property)**.** *Let A and B be events such that $A \cap B \neq \emptyset$ and $\overline{A} \cap B \neq \emptyset$, and let $\mathcal{X}$ be a non-empty finite $A \cap B$-consistent set of gambles, with $\{X, Y\} \subseteq \mathcal{X}$ such that $AX = AY$. Then $X \in \mathrm{opt}(\mathcal{X}|A \cap B)$ implies $Y \in \mathrm{opt}(\mathcal{X}|A \cap B)$ whenever there is a non-empty finite $\overline{A} \cap B$-consistent set of gambles $\mathcal{Z}$ such that, for at least one $Z \in \mathcal{Z}$,*

$$AX + \overline{A}Z \in \mathrm{opt}(A\mathcal{X} + \overline{A}\mathcal{Z}|B).$$

This property requires that, if two gambles agree on A, it is not possible for exactly one to be optimal, conditional on any subset of A, unless there is no suitable $\mathcal{Z}$.

Property 2 (Insensitivity of optimality to the omission of non-optimal elements)**.** *For any event* $A \neq \emptyset$, *and any non-empty finite A-consistent sets of gambles $\mathcal{X}$ and $\mathcal{Y}$,*

$$\mathrm{opt}(\mathcal{X}|A) \subseteq \mathcal{Y} \subseteq \mathcal{X} \Rightarrow \mathrm{opt}(\mathcal{Y}|A) = \mathrm{opt}(\mathcal{X}|A).$$

If opt satisfies this property, then removing non-optimal elements from a set does not affect whether or not each of the remaining elements is optimal. The property is called 'insensitivity to the omission of non-optimal elements' by De Cooman and Troffaes [5], and 'property ϵ' by Sen [32] who attributes this designation to Douglas Blair.

Property 3 (Preservation of non-optimality under the addition of elements)**.** *For any event* $A \neq \emptyset$, *and any non-empty finite A-consistent sets of gambles $\mathcal{X}$ and $\mathcal{Y}$,*

$$\mathcal{Y} \subseteq \mathcal{X} \Rightarrow \mathrm{opt}(\mathcal{Y}|A) \supseteq \mathrm{opt}(\mathcal{X}|A) \cap \mathcal{Y}.$$

This is called 'property α' by Sen [32], Axiom 7 by Luce and Raiffa [19, p. 288], and 'independence of irrelevant alternatives' by Radner and Marschak [24].[2] It states that any gamble that is non-optimal in a set of gambles $\mathcal{Y}$ is non-optimal in any set of gambles containing $\mathcal{Y}$.

Properties 2 and 3 are together equivalent to the well-known property of path independence, which can often be checked more conveniently.

Property 4. *A choice function* opt *is* path independent *if, for any non-empty event A, and for any finite family of non-empty finite A-consistent sets of gambles $\mathcal{X}_1, \ldots, \mathcal{X}_n$,*

$$\mathrm{opt}\left(\bigcup_{i=1}^{n} \mathcal{X}_i \middle| A \right) = \mathrm{opt}\left(\bigcup_{i=1}^{n} \mathrm{opt}(\mathcal{X}_i|A) \middle| A \right).$$

Path independence appears frequently in the social choice literature. Plott [23] gives a detailed investigation of path independence and its possible justifications. Path independence is also equivalent to Axiom 7$'$ of Luce and Raiffa [19, p. 289].

Lemma 14 (Sen [32, Proposition 19])**.** *A choice function* opt *satisfies Properties 2 and 3 if and only if* opt *satisfies Property 4.*

Property 5 (Backward mixture property)**.** *For any events A and B such that $B \cap A \neq \emptyset$ and $B \cap \overline{A} \neq \emptyset$, any $B \cap \overline{A}$-consistent gamble Z, and any non-empty finite $B \cap A$-consistent set of gambles $\mathcal{X}$,*

$$\mathrm{opt}\left(A\mathcal{X} + \overline{A}Z | B \right) \subseteq A\,\mathrm{opt}(\mathcal{X}|A \cap B) + \overline{A}Z.$$

Theorem 15 (Backward induction theorem)**.** *Let* opt *be any choice function. The following conditions are equivalent.*

(A) *For any consistent decision tree T, it holds that* $\mathrm{back}_{\mathrm{opt}}(T) = \mathrm{norm}_{\mathrm{opt}}(T)$.
(B) opt *satisfies Properties 1, 2, 3, and 5.*

5. Application to Coherent Lower Previsions

In this section we investigate which of the choice functions for coherent lower previsions satisfy the conditions of Theorem 15(B). Some of the results are based on proofs for more general choice functions that can be found in the Appendix.

---

[2]This is different from several other properties bearing the same name, such as that of Arrow's Impossibility Theorem. For further discussion, see Ray [27].

5.1. **Maximality.** Maximality is a strict partial order, so Properties 2 and 3 hold, by Proposition A.1.

Proposition 16. *Maximality satisfies Property 1.*

Proof. We prove a stronger result. Let A be a non-empty event, $\mathcal{X}$ be a non-empty finite set of A-consistent gambles, and $\{X, Y\} \subseteq \mathcal{X}$ with $AX = AY$. We show that, for any event B such that $A \cap B \neq \emptyset$, $X \in \mathrm{opt}_{>_{\underline{P}|\cdot}}(\mathcal{X}|A \cap B)$ implies $Y \in \mathrm{opt}_{>_{\underline{P}|\cdot}}(\mathcal{X}|A \cap B)$.

If $X \in \mathrm{opt}_{>_{\underline{P}|\cdot}}(\mathcal{X}|A \cap B)$, then for every $Z \in \mathcal{X}$, $\underline{P}(Z - X|A \cap B) \leq 0$. But $AX = AY$ implies $(A \cap B)X = (A \cap B)Y$, and, by Proposition 3(i), $\underline{P}(Z - X|A \cap B) = \underline{P}(Z - Y|A \cap B)$, and so it immediately follows that $Y \in \mathrm{opt}_{>_{\underline{P}|\cdot}}(\mathcal{X}|A \cap B)$. $\qquad\square$

Proposition 17. *Maximality satisfies Property 5.*

Proof. Consider events A and B such that $A \cap B \neq \emptyset$ and $\overline{A} \cap B \neq \emptyset$, a non-empty finite set of $A \cap B$-consistent gambles $\mathcal{X}$, and an $\overline{A} \cap B$-consistent gamble Z. To establish Property 5, it suffices to show that for any $Y \in \mathcal{X}$,

$$Y \notin \mathrm{opt}_{>_{\underline{P}|\cdot}}(\mathcal{X}|A \cap B) \implies AY + \overline{A}Z \notin \mathrm{opt}_{>_{\underline{P}|\cdot}}(A\mathcal{X} \oplus \overline{A}Z|B).$$

If $Y \notin \mathrm{opt}_{>_{\underline{P}|\cdot}}(\mathcal{X}|A \cap B)$ then there is an $X \in \mathcal{X}$ with $\underline{P}(X - Y|A \cap B) > 0$. The result follows if we show that $\underline{P}(AX + \overline{A}Z - (AY + \overline{A}Z)|B) > 0$. By Proposition 3(ii)–(v),

$$
\begin{aligned}
0 &= \underline{P}(A(X - Y - \underline{P}(X - Y|A \cap B))|B) \\
&\leq \underline{P}(A(X - Y)|B) + \overline{P}(-A\underline{P}(X - Y|A \cap B)|B) \\
&= \underline{P}(A(X - Y)|B) - \underline{P}(A\underline{P}(X - Y|A \cap B)|B) \\
&= \underline{P}(A(X - Y)|B) - \underline{P}(A|B)\underline{P}(X - Y|A \cap B)
\end{aligned}
$$

where we relied on $\underline{P}(X - Y|A \cap B) > 0$ in the last step. So, indeed

$$\underline{P}(A(X - Y)|B) \geq \underline{P}(A|B)\underline{P}(X - Y|A \cap B) > 0.$$

$\qquad\square$

Corollary 18. *For any consistent decision tree T, it holds that*

$$\mathrm{back}_{\mathrm{opt}_{>_{\underline{P}|\cdot}}}(T) = \mathrm{norm}_{\mathrm{opt}_{>_{\underline{P}|\cdot}}}(T).$$

Proof. Immediate, from Propositions A.1, 16, and 17, and Theorem 15. $\qquad\square$

5.2. **E-admissibility.** Since E-admissibility is a union of maximality choice functions we have:

Corollary 19. *For any consistent decision tree T, it holds that*

$$\mathrm{back}_{\mathrm{opt}_{\mathcal{M}}}(T) = \mathrm{norm}_{\mathrm{opt}_{\mathcal{M}}}(T).$$

Proof. Immediate, from Proposition A.2, Corollary 18, and Theorem 15. $\qquad\square$

Further, from Theorem A.3, we have:

Corollary 20. *For any consistent decision tree T,*

$$\mathrm{norm}_{\mathrm{opt}_{\mathcal{M}}}(T) = \mathrm{norm}_{\mathrm{opt}_{\mathcal{M}}}(\mathrm{back}_{\mathrm{opt}_{>_{\underline{P}|\cdot}}}(T)).$$

	A	$\overline{A}$
X	1	1
Y	1.5	3.5
Z	0	4

	$\underline{P}(\cdot|B)$	$\overline{P}(\cdot|B)$
X	1	1
Y	2	3
Z	1	3

	$\underline{P}$	$\overline{P}$
$BX + \overline{B}Z$	1	2
$BY + \overline{B}Z$	1.5	3

TABLE 2. Gambles and their lower and upper previsions for Example 21.

5.3. **Interval Dominance.** By Proposition A.1, interval dominance satisfies Properties 2 and 3, and it satisfies Property 1 because $AX = AY$ implies $\underline{P}(X|A) = \underline{P}(Y|A)$ and $\overline{P}(X|A) = \overline{P}(Y|A)$. We now show that interval dominance fails Property 5.

Example 21. *Suppose A and B are events, and X, Y, and Z are the gambles given in Table 2. Let $\mathcal{M}$ contain all mass functions P such that A and B are independent, $1/4 \leq P(A) \leq 3/4$, and $P(B) = 1/2$. Let $\underline{P}$ be the lower envelope of $\mathcal{M}$.*

Lower and upper previsions of relevant gambles are given in Table 2; for example,

$$\overline{P}(BY + \overline{B}Z) = \max_{P \in \mathcal{M}} P(BY + \overline{B}Z) = \max_{P \in \mathcal{M}} P(B)P(Y|B) + P(\overline{B})P(Z|\overline{B})$$

$$= \frac{1}{2} \max_{P \in \mathcal{M}}(P(Y) + P(Z)) = \frac{1}{2} \max_{p \in [\frac{1}{4}, \frac{3}{4}]} (1.5(1-p) + 3.5p + 4p) = 3$$

and similar for all other gambles. Clearly, Y interval dominates X conditional on B, however, $BY + \overline{B}Z$ does not interval dominate $BX + \overline{B}Z$, violating Property 5.

Even though interval dominance violates Property 5, it can still be of use in backward induction. It is easily shown that (see for instance Troffaes [34])

$$\mathrm{opt}_{>_{\underline{P}|\cdot}}(\mathcal{X}|A) \subseteq \mathrm{opt}_{\sqsupset_{\underline{P}|\cdot}}(\mathcal{X}|A).$$

By Theorem A.3, we therefore have:

Corollary 22. *For any consistent decision tree T,*

$$\mathrm{norm}_{\mathrm{opt}_{>_{\underline{P}|\cdot}}}(T) = \mathrm{norm}_{\mathrm{opt}_{>_{\underline{P}|\cdot}}}(\mathrm{back}_{\mathrm{opt}_{\sqsupset_{\underline{P}|\cdot}}}(T))$$

$$\mathrm{norm}_{\mathrm{opt}_{\mathcal{M}}}(T) = \mathrm{norm}_{\mathrm{opt}_{\mathcal{M}}}(\mathrm{back}_{\mathrm{opt}_{\sqsupset_{\underline{P}|\cdot}}}(T))$$

It can also be shown that $\mathrm{back}_{\mathrm{opt}_{\sqsupset_{\underline{P}|\cdot}}}(T) \subseteq \mathrm{norm}_{\mathrm{opt}_{\sqsupset_{\underline{P}|\cdot}}}(T)$ for all T, so all strategies found by backward induction will be optimal with respect to $\mathrm{opt}_{\sqsupset_{\underline{P}|\cdot}}$.

5.4. **Γ-maximin.** Γ-maximin fails Theorem 15(A): see for example Seidenfeld [30, Sequential Example 1, pp. 75–77]. Since Γ-maximin is induced by an ordering, it satisfies Properties 2 and 3 by Proposition A.1. As for interval dominance, Γ-maximin satisfies Property 1. Hence, Γ-maximin must fail Property 5. Indeed, backward induction can fail in a particularly serious way: it can select a single gamble that is inferior to another normal form gamble. Hence, backward induction may not find any Γ-maximin gambles.

6. THE OIL WILDCATTER EXAMPLE

We now illustrate our algorithm using the same example as Kikuti et al. [14, Fig. 2]. Fig. 4 depicts the decision tree, with utiles in units of $10000. The subject must decide whether to drill for oil (d_2) or not (d_1). Drilling costs 7 and provides a return of 0, 12, or 27 depending on the richness of the site. The events S_1 to S_3 represent the different yields, with S_1 being the least profitable and S_3 the most. The subject may pay 1 to test the site before deciding whether to
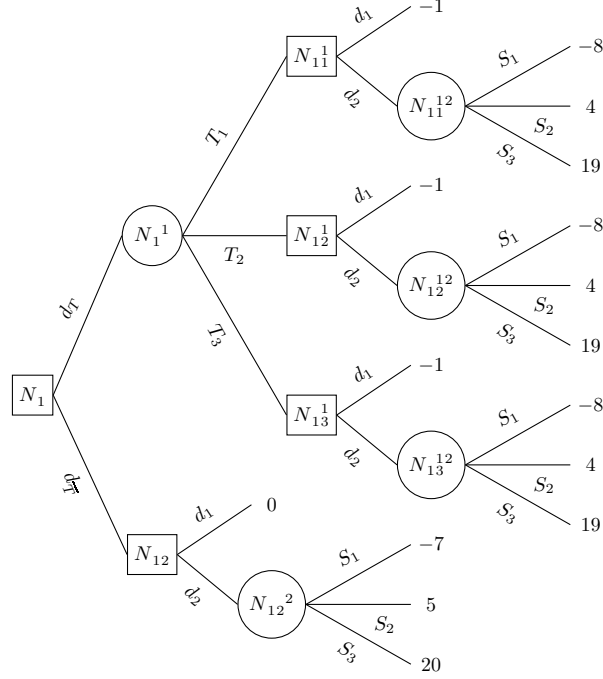
FIGURE 4. Decision tree for the oil wildcatter.

T_1	T_2	T_3
0.183, 0.222	0.333, 0.363	0.444, 0.454

TABLE 3. Unconditional lower and upper probabilities $\underline{P}(T_i)$ and $\overline{P}(T_i)$ for oil example.

	T_1	T_2	T_3
S_1	0.547, 0.653	0.222, 0.333	0.111, 0.166
S_2	0.222, 0.272	0.363, 0.444	0.333, 0.363
S_3	0.125, 0.181	0.250, 0.363	0.471, 0.556

TABLE 4. Conditional lower and upper probabilities $\underline{P}(S_i|T_i)$ and $\overline{P}(S_i|T_i)$ for oil example.

drill; this gives one of three results T_1 to T_3, where T_1 is the most pessimistic and T_3 the most optimistic.

Lower and upper probabilities are given for each T_i (Table 3), and for each S_i conditional on T_i (Table 4). (Some intervals are tighter than those in Kikuti et al., since their values are incoherent—we corrected these by natural extension [35, §3.1].)

By *marginal extension* [35, §6.7.2], the lower prevision of a gamble Z is then

$$\underline{P}(Z) = \underline{P}(T_1\underline{P}(Z|T_1) + T_2\underline{P}(Z|T_2) + T_3\underline{P}(Z|T_3)).$$

Let $X = -7S_1+5S_2+20S_3$, and again let $T_*^* = \mathrm{st}_{N_*^*}(T)$. Since we will only be concerned with maximality, and normal form decisions in this problem are uniquely identified by their gambles,
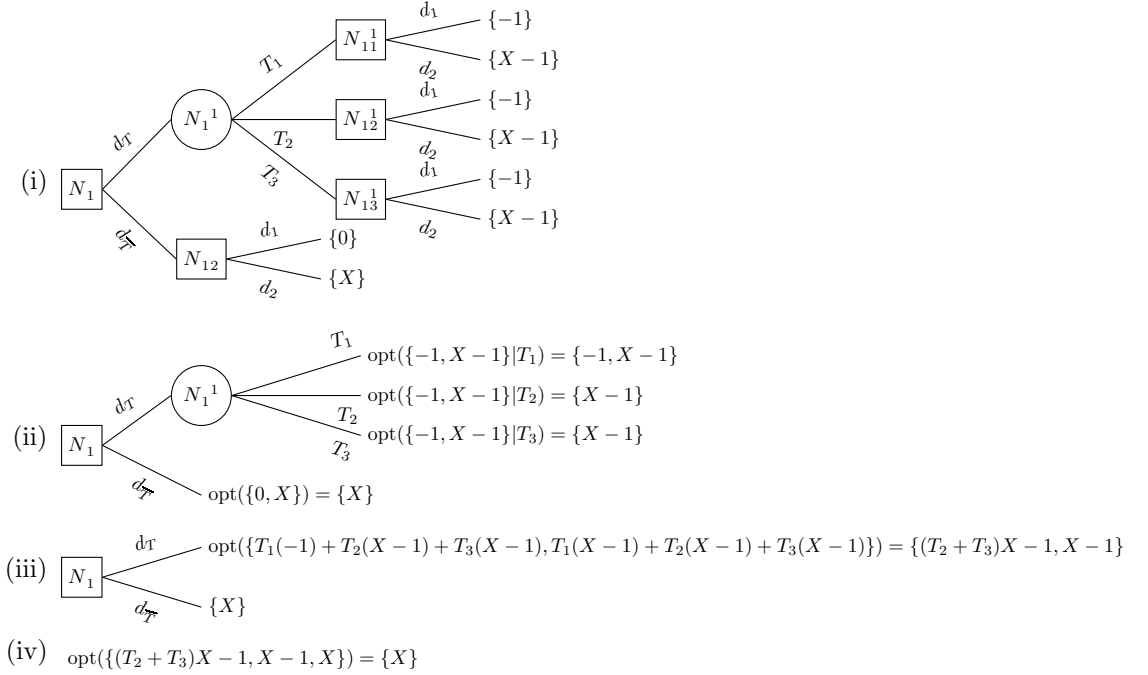
FIGURE 5. Solving the oil wildcatter example by normal form backward induction.

we can conveniently work with gambles in this example. Therefore, we use the following notation:

$$\mathrm{opt} = \mathrm{opt}_{>_{\underline{P}|.}} \qquad \mathrm{back} = \mathrm{gamb} \circ \mathrm{back}_{\mathrm{opt}_{>_{\underline{P}|.}}} \qquad \mathrm{norm} = \mathrm{gamb} \circ \mathrm{norm}_{\mathrm{opt}_{>_{\underline{P}|.}}}$$

Fig. 5 depicts the process of backward induction described next.

(i) Of course, $\mathrm{back}(\cdot)$ at the final chance nodes simply reports the gamble: $\mathrm{back}(T_{11}^{12}) = \mathrm{back}(T_{11}^{22}) = \mathrm{back}(T_{11}^{32}) = \{X - 1\}$, and $\mathrm{back}(T_{12}^{2}) = \{X\}$.

(ii) For T_{11}^{1}, we must find $\underline{P}((X - 1) - (-1)|T_1)$ and $\underline{P}(-1 - (X - 1)|T_1)$. These lower previsions can be computed using Table 4 as follows: X will have lowest expected value when the worst outcome S_1 is most likely (probability 0.653) and the best outcome S_3 is least likely (probability 0.125), and so the probability of S_2 is 0.222. So, $\underline{P}(X|T_1) = -7 \times 0.653 + 5 \times 0.222 + 20 \times 0.125 = -0.961$. Similarly, $\underline{P}(-X|T_1) = -1.151$. Neither of these is positive, so $\mathrm{back}(T_{11}^{1}) = \{X - 1, -1\}$.

For T_{12}^{1}, $\underline{P}((X-1)-(-1)|T_2) = 4.754$, and therefore d_2 dominates d_1, so $\mathrm{back}(T_{12}^{1}) = \{X - 1\}$. Similarly, $\underline{P}((X-1)-(-1)|T_3) = 10.073$, so $\mathrm{back}(T_{13}^{1}) = \{X - 1\}$.

For T_{12}, we need to find $\underline{P}(X - 0)$. By marginal extension we have

$$\begin{aligned}
\underline{P}(X) &= \underline{P}(T_1\underline{P}(X|T_1) + T_2\underline{P}(X|T_2) + T_3\underline{P}(X|T_3)) \\
&= \underline{P}(-0.961T_1 + 4.754T_2 + 10.073T_3) \\
&= 0.222 \times -0.961 + 0.334 \times 4.754 + 0.444 \times 10.073 = 5.846906.
\end{aligned}$$

This is greater than zero, so $\mathrm{back}(T_{12}) = \{X\}$.

(iii) At T_1^{1} there are two potentially optimal gambles: $T_1(X - 1) + T_2(X - 1) + T_3(X - 1) = X - 1$ and $T_1(-1) + T_2(X - 1) + T_3(X - 1) = (T_2 + T_3)X - 1$. We must find $\underline{P}((X-1) - ((T_2+T_3)X - 1)) = \underline{P}(T_1 X)$ and $\underline{P}(((T_2+T_3)X - 1) - (X-1)) = \underline{P}(-T_1 X)$.

Using marginal extension,

$$\underline{P}(T_1 X) = \underline{P}(T_1 \underline{P}(X|T_1)) = \underline{P}(-0.961 T_1) = 0.222 \times -0.961 = -0.213342 < 0,$$
$$\underline{P}(-X T_1) = \underline{P}(T_1 \underline{P}(-X|T_1)) = \underline{P}(-1.151 T_1) = 0.222 \times -1.151 = -0.255522 < 0,$$

so $\text{back}(T_1{}^1) = \{X - 1, (T_2 + T_3)X - 1\}$.

(iv) Finally, for T, we must consider $\{X, X - 1, (T_2 + T_3)X - 1\}$. It is clear that $\underline{P}(X - (X - 1)) = 1 > 0$, so $X - 1$ can be eliminated. It is also clear that if a gamble does not dominate $X - 1$ then is also does not dominate X, so by our calculation at $T_1{}^1$ we know that X is maximal. We finally have

$$\underline{P}(X - ((T_2 + T_3)X - 1)) = \underline{P}(T_1 X + 1) = \underline{P}(T_1 X) + 1 = -0.213342 + 1 > 0,$$

so $\text{back}(T) = \{X\}$. So, the optimal strategy is: do not test and just drill.

We found a single maximal strategy. By Corollary 20, it is also the unique E-admissible strategy. (Our solution differs from Kikuti et al. [14]; since they do not detail their calculations, we could not identify why.) Of course, if the imprecision was larger, we would have found more, but it does show that non-trivial sequential problems can give unique solutions even when probabilities are imprecise.

In this example, the usual normal form method requires comparing 10 gambles at once. By normal form backward induction, we only had to compare 2 gambles at once at each stage (except at the end, where we had 3), leading us much quicker to the solution: the computational benefit of normal form backward induction is obvious.

7. Conclusion

When solving sequential decision problems with limited knowledge, it may be impossible to assign a probability to every event. An alternative is to use a coherent lower prevision, or, equivalently, a closed convex set of probability mass functions. Under such a model there are several plausible generalizations of maximizing expected utility.

Given any criterion, we considered two methods of solving decision trees: the usual normal form, where the subject applies the criterion to the lists all strategies, and normal form backward induction, adapted from [14]. If they coincide, backward induction helps efficiently solving trees. If they differ, doubt is cast on the criterion's suitability.

In Theorem 15 we identified when the two methods coincide. We then applied these results to the choice functions for coherent lower previsions. As was already known, Γ-maximin fails Property 5. Interval dominance fails the same condition. However, and perhaps surprisingly, maximality and E-admissibility satisfy all conditions, in the case where lower probabilities are non-zero. If any lower probabilities are zero, then Property 5 fails (unsurprisingly, as in this case it already fails with precise probabilities).

When analysing choice functions, whether for lower previsions or not, usually Property 5 is most troublesome. Failing Property 1 would involve an unnatural form of conditioning, and path independence (Property 4) is a very natural consistency condition that one usually wants to satisfy before even considering decision trees.

We have not argued that a normal form operator, and norm$_{\text{opt}}$ in particular, gives the best solution to a decision tree. A normal form solution requires a policy to be specified and adhered to. The subject does this policy only by her own resolution: she of course can change her policy upon reaching a decision node [29]. One might argue that a normal form solution is only acceptable when the subject cannot change her mind (for example, if she instructs, in advance, others to carry out the actions).

Further, many choice functions cause norm$_{\text{opt}}$ to have undesirable properties, even when they satisfy Properties 1, 2, 3, and 5. For example, using maximality or E-admissibility with a lower

prevision allows the subject to choose to pay for apparently irrelevant information instead of making an immediate decision [31], and a gamble that is optimal in a subtree may become non-optimal in the full tree [8, 10, 11].

Moreover, normal form backward induction does not always help with computations. The need to store, at every stage, all optimal gambles, could be a burden. Secondly, if imprecision is large, causing only few gambles to be deleted, the set of optimal gambles at each stage will still eventually become too large. In such situations, a form of approximation may be necessary. Even so, we have shown that, perhaps surprisingly, the normal form *can* be solved exactly with backward induction, and when either trees or imprecision are not too large, the method will be computationally feasible.

References

[1] T. Augustin, *On decision making under ambiguous prior and sampling information*, ISIPTA 01: Proceedings of the Second International Symposium on Imprecise Probabilities (G. de Cooman, T.L. Fine, S. Moral, and T. Seidenfeld, eds.), 2001.

[2] George Boole, *An investigation of the laws of thought on which are founded the mathematical theories of logic and probabilities*, Walton and Maberly, London, 1854.

[3] Robert T. Clemen and Terence Reilly, *Making hard decisions*, Duxbury, Belmont CA, 2001.

[4] Marquis de Condorcet, *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*, L'Imprimerie Royale, Paris, 1785.

[5] G. De Cooman and Matthias C. M. Troffaes, *Dynamic programming for deterministic discrete-time systems with uncertain gain*, International Journal of Approximate Reasoning **39** (2005), no. 2-3, 257–278, `doi:10.1016/j.ijar.2004.10.004`.

[6] Bruno de Finetti, *Theory of probability: A critical introductory treatment*, Wiley, New York, 1974–5, Two volumes.

[7] Jonathan Gross and Jay Yellen, *Graph theory and its applications*, CRC Press, London, 1999.

[8] P. Hammond, *Consequentialist foundations for expected utility*, Theory and Decision **25** (1988), no. 1, 25–78.

[9] David Harmanec, *A generalization of the concept of markov decision process to imprecise probabilities*, 1st International Symposium on Imprecise Probabilities and Their Applications, 1999, pp. 175–182.

[10] N. Huntley and Matthias C. M. Troffaes, *Characterizing factuality in normal form sequential decision making*, ISIPTA'09: Proceedings of the Sixth International Symposium on Imprecise Probability: Theories and Applications (Thomas Augustin, Frank P. A. Coolen, Serafin Moral, and Matthias C. M. Troffaes, eds.), Jul 2009, pp. 239–248.

[11] ———, *Subtree perfectness, backward induction, and normal-extensive form equivalence for single agent sequential decision making under arbitrary choice functions*, arXiv:1109.3607v1 [math.ST], September 2011, `arXiv:1109.3607`.

[12] Nathan Huntley and Matthias C. M. Troffaes, *An efficient normal form solution to decision trees with lower previsions*, Soft Methods for Handling Variability and Imprecision (Berlin) (Didier Dubois, M. Asunción Lubiano, Henri Prade, María Ángeles Gil, Przemyslaw Grzegorzewski, and Olgierd Hryniewicz, eds.), Advances in Soft Computing, Springer, Sep 2008, pp. 419–426.

[13] J. Jaffray, *Rational decision making with imprecise probabilities*, 1st International Symposium on Imprecise Probabilities and Their Applications, 1999, pp. 183–188.

[14] D. Kikuti, F. Cozman, and C.P. de Campos, *Partially ordered preferences in decision trees: Computing strategies with imprecision in probabilities*, IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling (R. Brafman and U. Junker, eds.), 2005, pp. 118–123.

[15] H.E. Kyburg, *Rational belief*, Behavioral and Brain Sciences **8** (1983), no. 2, 231–273.

[16] I. LaValle and K. Wapman, *Rolling back decision trees requires the independence axiom!*, Management Science **32** (1986), no. 3, 382–385.

[17] I. Levi, *The enterprise of knowledge*, MIT Press, London, 1980.

[18] D. V. Lindley, *Making decisions*, 2nd ed., Wiley, London, 1985.

[19] R.D. Luce and H. Raiffa, *Games and decisions: introduction and critical survery*, Wiley, 1957.

[20] M.J. Machina, *Dynamic consistency and non-expected utility models of choice under uncertainty*, Journal of Economic Literature **27** (1989), 1622–1688.

[21] E. F. McClennen, *Rationality and dynamic choice: Foundational explorations*, Cambridge University Press, Cambridge, 1990.

[22] Enrique Miranda, *A survey of the theory of coherent lower previsions*, International Journal of Approximate Reasoning **48** (2008), no. 2, 628–658.

[23] C.R. Plott, *Path independence, rationality, and social choice*, Econometrica **41** (1973), no. 6, 1075–1091.

[24] R. Radner and J. Marschak, *Note on some proposed decision criteria*, Decision Processes (R. M. Thrall, C. H Coombs, and R. L. Davies, eds.), John Wiley, New York, 1954, pp. 61–68.

[25] H. Raiffa and R. Schlaifer, *Applied statistical decision theory*, Harvard University Press, Boston, 1961.

[26] Frank P. Ramsey, *Truth and probability*, Foundations of Mathematics and other Logical Essays (R. B. Braithwaite, ed.), Routledge and Kegan Paul, London, 1931, Published posthumously, pp. 156–198.

[27] P. Ray, *Independence of irrelevant alternatives*, Econometrica **41** (1973), no. 5, 987–991.

[28] Sheldon Ross, *A first course in probability*, 7th ed., Pearson Prentice Hall, 2006.

[29] T. Seidenfeld, *Decision theory without 'independence' or without 'ordering': What is the difference?*, Economics and Philosophy **4** (1988), 267–290.

[30] ———, *A contrast between two decision rules for use with (convex) sets of probabilities: Γ-maximin versus E-admissibility*, Synthese **140** (2004), 69–88.

[31] T. Seidenfeld, M. J. Schervish, and J. B Kadane, *Coherent choice functions under uncertainty*, 5th International Symposium on Imprecise Probability: Theories and Applications, 2007, pp. 385–394.

[32] A. K. Sen, *Social choice theory: A re-examination*, Econometrica **45** (1977), no. 1, 53–89.

[33] Cedric A. B. Smith, *Consistency in statistical inference and decision*, Journal of the Royal Statistical Society **B** (1961), no. 23, 1–37.

[34] Matthias C. M. Troffaes, *Decision making under uncertainty using imprecise probabilities*, International Journal of Approximate Reasoning **45** (2007), 17–29.

[35] P. Walley, *Statistical reasoning with imprecise probabilities*, Chapman and Hall, London, 1991.

[36] Peter M. Williams, *Notes on conditional previsions*, Tech. report, School of Math. and Phys. Sci., Univ. of Sussex, 1975.

[37] ———, *Notes on conditional previsions*, International Journal of Approximate Reasoning **44** (2007), no. 3, 366–383.

[38] Marco Zaffalon, Keith Wesnes, and Orlando Petrini, *Reliable diagnoses of dementia by the naive credal classifier inferred from incomplete cognitive data*, Artificial Intelligence in Medicine **29** (2003), no. 1–2, 61–79.

APPENDIX A. RESULTS FOR GENERAL CHOICE FUNCTIONS

This appendix details intermediate results required for the proofs in Section 3. Since the results are applicable for choice functions that are nothing to do with coherent lower previsions, and so may be useful for investigating other uncertainty models, we present them separately.

Proposition A.1. *For each non-empty event A, let $\succ_A$ be any strict partial order on A-consistent gambles. The choice function induced by these strict partial orders, that is,*

$$\mathrm{opt}_{\succ|.}(\mathcal{X}|A) = \{X \in \mathcal{X} \colon (\forall Y \in \mathcal{X})(Y \not\succ_A X)\}$$

satisfies Properties 2 and 3.

Proof. By Lemma 14, it suffices to show that $\mathrm{opt}_{\succ|.}$ is path independent. Let $\mathcal{X}_1$, $\ldots$, $\mathcal{X}_n$ be non-empty finite sets of A-consistent gambles, and let A be a non-empty event. Let $\mathcal{X} = \bigcup_{i=1}^{n} \mathcal{X}_i$ and $\mathcal{Z} = \bigcup_{i=1}^{n} \mathrm{opt}_{\succ|.}(\mathcal{X}_i|A)$. We show must show that

$$\mathrm{opt}_{\succ|.}(\mathcal{X}|A) = \mathrm{opt}_{\succ|.}(\mathcal{Z}|A). \tag{3}$$

By definition,

$$\mathrm{opt}_{\succ|.}(\mathcal{Z}|A) = \{Z \in \mathcal{Z} \colon (\forall Y \in \mathcal{Z})(Y \not\succ_A Z)\},$$

and, observe that, if $X \in \mathcal{X}$ but $X \notin \mathcal{Z}$, by transitivity of $\succ_A$ and finiteness of $\mathcal{X}$, there is a $Y \in \mathcal{Z}$ such that $Y \succ_A X$. Therefore again by transitivity of $\succ_A$, for any $Z \in \mathcal{Z}$ such that $X \succ_A Z$, we have $Y \succ_A Z$. So,

$$= \{Z \in \mathcal{Z} \colon (\forall Y \in \mathcal{X})(Y \not\succ_A Z)\},$$

and once again, by definition of $\mathcal{Z}$, if $X \in \mathcal{X}$ but $X \notin \mathcal{Z}$ there is a $Y \in \mathcal{X}$ such that $Y \succ_A X$, so we have

$$= \{X \in \mathcal{X} \colon (\forall Y \in \mathcal{X})(Y \not\succ_A X)\}$$
$$= \mathrm{opt}_{\succ|.}(\mathcal{X}|A).$$

$\square$

Proposition A.2. *Let $\{\mathrm{opt}_i \colon i \in \mathcal{I}\}$ be a family of choice functions. For any non-empty event A and any non-empty finite set of A-consistent gambles $\mathcal{X}$, let*

$$\mathrm{opt}(\mathcal{X}|A) = \bigcup_{i \in \mathcal{I}} \mathrm{opt}_i(\mathcal{X}|A).$$

(i) If each opt_i satisfies Property 2, then so does opt.
(ii) If each opt_i satisfies Property 3, then so does opt.
(iii) If each opt_i satisfies Property 5, then so does opt.
(iv) If each opt_i satisfies Properties 1, 3, and 5, then opt satisfies Property 1.

Proof. (i). By definition of opt and by assumption, for any finite non-empty sets of gambles $\mathcal{X}$ and $\mathcal{Y}$ such that $\mathcal{Y} \subseteq \mathcal{X}$ and for any $i \in \mathcal{I}$, $\mathrm{opt}_i(\mathcal{X}|A) \subseteq \mathrm{opt}(\mathcal{X}|A) \subseteq \mathcal{Y}$, and therefore by Property 2, $\mathrm{opt}_i(\mathcal{X}|A) = \mathrm{opt}_i(\mathcal{Y}|A)$. Whence,

$$\mathrm{opt}(\mathcal{Y}|A) = \bigcup_{i \in \mathcal{I}} \mathrm{opt}_i(\mathcal{Y}|A) = \bigcup_{i \in \mathcal{I}} \mathrm{opt}_i(\mathcal{X}|A) = \mathrm{opt}(\mathcal{X}|A).$$

(ii). By assumption, for any finite non-empty sets of gambles $\mathcal{X}$ and $\mathcal{Y}$ such that $\mathcal{Y} \subseteq \mathcal{X}$ and for any $i \in \mathcal{I}$, $\mathrm{opt}_i(\mathcal{Y}|A) \supseteq \mathrm{opt}_i(\mathcal{X}|A) \cap \mathcal{Y}$. Therefore,

$$\mathrm{opt}(\mathcal{Y}|A) = \bigcup_{i \in \mathcal{I}} \mathrm{opt}_i(\mathcal{Y}|A) \supseteq \bigcup_{i \in \mathcal{I}} (\mathrm{opt}_i(\mathcal{X}|A) \cap \mathcal{Y})$$
$$= \mathcal{Y} \cap \bigcup_{i \in \mathcal{I}} \mathrm{opt}_i(\mathcal{X}|A) = \mathrm{opt}(\mathcal{X}|A) \cap \mathcal{Y}.$$

(iii). By assumption, for any non-empty finite set of gambles $\mathcal{X}$, any gamble Z, any events A and B such that $A \cap B \neq \emptyset$, and for any $i \in \mathcal{I}$,

$$\mathrm{opt}_i(A\mathcal{X} \oplus \overline{A}Z|B) \subseteq A\,\mathrm{opt}_i(\mathcal{X}|A \cap B) \oplus \overline{A}Z,$$

whence

$$\mathrm{opt}(A\mathcal{X} \oplus \overline{A}Z|B) = \bigcup_{i \in \mathcal{I}} \mathrm{opt}_i(A\mathcal{X} \oplus \overline{A}Z|B)$$
$$\subseteq \bigcup_{i \in \mathcal{I}} (A\,\mathrm{opt}_i(\mathcal{X}|A \cap B) \oplus \overline{A}Z)$$
$$= \overline{A}Z \oplus A \bigcup_{i \in \mathcal{I}} \mathrm{opt}_i(\mathcal{X}|A \cap B)$$
$$= \overline{A}Z \oplus A\,\mathrm{opt}(\mathcal{X}|B).$$

(iv). Let A and B be events such that $A \cap B \neq \emptyset$ and $\overline{A} \cap B \neq \emptyset$, $\mathcal{Z}$ be a non-empty finite set of $\overline{A} \cap B$-consistent gambles, and $\mathcal{X}$ be a non-empty finite set of $A \cap B$-consistent gambles such that there is $\{X, Y\} \subseteq \mathcal{X}$ with $AX = AY$. Suppose that there is a $Z \in \mathcal{Z}$ such that $AX + \overline{A}Z \in \mathrm{opt}(A\mathcal{X} + \overline{A}\mathcal{Z}|B)$.

By definition of opt, there is a j such that $AX + \overline{A}Z \in \mathrm{opt}_j(A\mathcal{X} + \overline{A}\mathcal{Z}|B)$. We show that both X and Y are in $\mathrm{opt}_j(\mathcal{X}|A \cap B)$, and therefore are both in $\mathrm{opt}(\mathcal{X}|A \cap B)$. It follows from Properties 3 and 5 that

$$\mathrm{opt}_j(A\mathcal{X} + \overline{A}\mathcal{Z}|B) \subseteq A\,\mathrm{opt}_j(\mathcal{X}|A \cap B) + \overline{A}\,\mathrm{opt}_j(\mathcal{Z}|\overline{A} \cap B).$$

Therefore, there is a $V \in \mathrm{opt}_j(\mathcal{X}|A \cap B)$ with $AV = AX$. Finally, opt_j satisfies Property 1, and therefore both X and Y must be in $\mathrm{opt}_j(\mathcal{X}|A \cap B)$. This establishes Property 1 for opt. $\qquad \square$

The final result is the following: if opt_1 satisfies the necessary properties, opt_2 does not, but $\mathrm{opt}_1 \subseteq \mathrm{opt}_2$, then we can use $\mathrm{opt}_1(\mathrm{back}_{\mathrm{opt}_2}(\cdot))$ to find $\mathrm{norm}_{\mathrm{opt}_1}$. This could be of interest in situations where opt_2 is much more computationally efficient than opt_1, and still eliminates enough gambles to be useful.

Theorem A.3. *Let* opt_1 *and* opt_2 *be choice functions such that* opt_1 *satisfies Properties 1, 2, 3, and 5, and for any non-empty event A and any non-empty finite set of A-consistent gambles* $\mathcal{X}$,

$$\mathrm{opt}_1(\mathcal{X}|A) \subseteq \mathrm{opt}_2(\mathcal{X}|A).$$

Then, for any consistent decision tree T,

$$(4) \qquad\qquad \mathrm{norm}_{\mathrm{opt}_1}(T) = \mathrm{norm}_{\mathrm{opt}_1}(\mathrm{back}_{\mathrm{opt}_2}(T)).$$

Durham University, Department of Mathematical Sciences, Science Laboratories, South Road, Durham DH1 3LE, United Kingdom
E-mail address: nathan.huntley@durham.ac.uk

Durham University, Department of Mathematical Sciences, Science Laboratories, South Road, Durham DH1 3LE, United Kingdom
E-mail address: matthias.troffaes@gmail.com